FUNDAMENTAL SCIENTIFIC REPORTS IN MULTIDISCIPLINARY AREAS

Vol 2 Issue 03 (**2026**)

www.fsrma.org

ONLINE ISSN: 3134-7517

PRINT ISSN: 3134-7509

Volatility-Surface Encoders for Regime-Robust Realized Volatility Forecasting Models

Lok-Yin Ng, Jack Tse

Department of Economics and Finance, Faculty of Commerce, Hong Kong Shue Yan University, Hong Kong, Hong Kong SAR, China

Abstract.

The accurate forecasting of realized volatility is a fundamental task in quantitative finance, with extensive applications in derivative pricing, portfolio optimization, and risk management. Traditional econometric models and standard machine learning approaches frequently fail to maintain predictive accuracy during abrupt market regime shifts, such as transitions from calm periods to financial crises. This paper introduces a novel deep learning framework that incorporates a Volatility-Surface Encoder to achieve regime-robust forecasting of realized volatility. By treating the entire implied volatility surface as a two-dimensional spatial input, the proposed model employs convolutional autoencoder architectures to extract high-dimensional, forward-looking latent representations of market expectations. These spatial representations are subsequently integrated with historical high-frequency realized volatility measurements using temporal neural networks. The comprehensive empirical analysis utilizes a large-scale dataset of equity index options and high-frequency asset returns, evaluating performance across distinct market regimes. The findings indicate that the inclusion of spatially encoded implied volatility features significantly reduces out-of-sample forecasting errors, particularly during periods characterized by structural breaks and elevated market stress. The proposed methodology provides a robust alternative to purely backward-looking autoregressive models, demonstrating the profound informational value embedded in the structural dynamics of option pricing surfaces.

Keywords: *Realized Volatility, Implied Volatility Surface, Deep Learning, Regime Robustness, Forecasting Models, Representation Learning*

1. Introduction

1.1 The Challenge of Volatility Forecasting

The precise estimation and projection of financial market volatility remain central pillars of modern quantitative finance, underpinning the practical applications of risk management, derivative valuation, and dynamic portfolio optimization. Realized volatility, constructed from high-frequency intraday asset returns, provides a highly informative non-parametric measure of ex-post variation. Arbitrage pricing provides a factor-based financial reference, useful for distinguishing systematic exposures from unexplained residual variation [1]. However, projecting this measure into the future continues to challenge both traditional econometric models and contemporary machine learning algorithms. The primary difficulty stems from the complex, non-

linear dynamics inherent in financial markets and the pervasive presence of latent market regimes and structural breaks. During periods of relative economic tranquility, volatility dynamics exhibit strong persistence, slow mean reversion, and highly predictable autoregressive patterns. Conversely, during financial crises, exogenous macroeconomic shocks, or sudden liquidity events, volatility spikes dramatically, and the historical correlation structures between asset classes rapidly break down. Standard linear frameworks typically fail to adapt rapidly to these sudden regime shifts because their parameters are estimated over long historical rolling windows that inadvertently smooth out abrupt transitions and ignore the forward-looking signals generated by market participants. Zhao and colleagues examine dynamic AI-Driven architecture for real estate asset tokenization on blockchain [2]. The reliance on purely backward-looking historical data is a significant limitation in standard forecasting models. Financial markets are fundamentally driven by the aggregated expectations of investors, which are not always perfectly reflected in the historical sequence of realized returns. When a market transitions from a low-volatility regime to a high-volatility regime, historical data becomes momentarily uninformative, leading to substantial forecasting lag and significant financial losses for risk managers relying on these models. Consequently, there is a critical need in the academic literature and industry practice for forecasting methodologies that can anticipate regime shifts or adapt to them instantaneously by incorporating forward-looking informational variables. This requirement motivates the exploration of alternative data sources and advanced computational techniques capable of synthesizing complex, high-dimensional market signals into actionable forecasts. Jin and colleagues examine evidence froma52-country panel analysis [3].

1.2 Motivation for Volatility Surface Encoders

Options markets provide a profound, forward-looking perspective on asset price dynamics through the mechanism of implied volatility. Implied volatility, derived by inverting standard option pricing formulas, represents the market consensus of future volatility over the life of the option contract. When mapped across various strike prices and times to maturity, these individual implied volatilities form a complex, three-dimensional structure known as the implied volatility surface. The precise topology of this surface, including the steepness of the volatility skew and the curvature of the volatility smile, encapsulates critical information regarding market sentiment, tail risk expectations, and the probability of sudden asset price drops. Despite its rich informational content, the full implied volatility surface is rarely utilized in standard realized volatility forecasting models due to its high dimensionality, persistent multicollinearity, and missing data problems inherent in illiquid option contracts. Another relevant line of investigation concerns finance-Structured temporal transformers for realized volatility forecasting from implied volatility surfaces [4]. To overcome these structural barriers, this study proposes the implementation of Volatility-Surface Encoders. These encoders leverage representation learning architectures, specifically convolutional autoencoders, to compress the high-dimensional grid of the implied volatility surface into a compact, dense latent vector. This latent representation captures the fundamental geometric properties of the surface without succumbing to the noise of individual option pricing anomalies. The Black-Scholes option-pricing formulation provides a classical benchmark whose distributional assumptions should remain explicit when evaluating a more flexible uncertainty model [5]. By integrating this forward-looking latent vector with historical realized volatility sequences, the resulting forecasting model bridges the gap between historical realization and market expectation. This synthesis is particularly vital for achieving regime robustness, as the options market often prices in the probability of regime transitions long before they manifest in the underlying asset returns. The integration of spatial feature extraction from

option panels with temporal sequence modeling provides a holistic framework for predicting volatility across diverse and highly volatile market environments. For a broader comparison, related research addresses comprehensive survey of Cross-Modal retrieval: breaking through modal barriers [6].

1.3 Structure of the Research

This paper comprehensively details the development, optimization, and empirical validation of the Volatility-Surface Encoder framework. The subsequent sections are organized to provide a systematic exposition of the theoretical foundations and the methodological innovations. Section 2 presents an extensive review of the pertinent literature, covering the evolution of realized volatility models, the informational content of implied volatility surfaces, and the application of deep learning techniques in financial econometrics. Section 3 delineates the methodology and model architecture, providing a detailed explanation of the data transformation processes, the structural design of the convolutional autoencoders, and the formulation of the integrated temporal forecasting network. Section 4 presents the empirical results, offering a rigorous performance analysis of the proposed models against established benchmarks, with a specific focus on evaluating predictive accuracy across disparate market regimes characterized by varying levels of stress. Finally, Section 5 concludes the research by summarizing the primary contributions, discussing the practical implications for quantitative trading and risk management, and suggesting promising avenues for future academic inquiry. Another relevant line of investigation concerns comparative explanations via counterfactual reasoning in recommendations [7].

2. Literature Review and Theoretical Background

2.1 Evolution of Realized Volatility Models

The academic pursuit of modeling and forecasting asset price volatility has undergone several paradigm shifts over the past four decades. Early approaches predominantly relied on parametric models applied to low-frequency daily returns, most notably the Generalized Autoregressive Conditional Heteroskedasticity models. The cross-section of expected stock returns provides an empirical factor-model reference for evaluating whether a learned signal adds information beyond established exposures [8]. While mathematically elegant, these models suffer from delayed responsiveness to market shocks due to their dependence on single daily observations. The advent of high-frequency trading data facilitated the transition toward non-parametric measures, specifically realized volatility, which aggregates intraday squared returns to provide a highly accurate estimate of integrated variance. The foundational properties of realized volatility, including its long-memory characteristics and approximate log-normal distribution, led to the development of specialized forecasting frameworks. Another relevant line of investigation concerns token design for favor trading with dynamic membership [9]. The most prominent among these frameworks is the Heterogeneous Autoregressive model of Realized Volatility. This model successfully captures the long-memory property of volatility by regressing future volatility on daily, weekly, and monthly averages of past realized volatility. The simplicity and robust empirical performance of this model have established it as the dominant benchmark in the literature. However, subsequent research has highlighted its vulnerability to structural breaks and sudden regime shifts. During the global financial crisis, standard autoregressive models exhibited severe under-prediction of volatility spikes and persistent over-prediction during the subsequent market recovery. To address these deficiencies, extensions have been proposed that incorporate jump components, leverage effects, and macroeconomic variables. Despite these enhancements, models restricted to historical price data fundamentally lack the anticipatory capabilities required to navigate abrupt transitions in market state. Time to productive use is considered in the workforce-

agility study [10]. Deployment of a volatility encoder has a comparable operational dimension: analysts may need time to learn its failure modes, so usability and model forecast error should be assessed independently.

2.2 Information Content of Implied Volatility Surfaces

The limitations of purely backward-looking models have directed researchers toward the options market, which naturally aggregates forward-looking expectations. Implied volatility is widely acknowledged as an unbiased, albeit risk-premium adjusted, predictor of future realized volatility. Early studies focused primarily on at-the-money implied volatility indices, demonstrating that the inclusion of these single-point metrics significantly improved the predictive accuracy of autoregressive models. However, reducing the entire options market to a single index discards massive amounts of cross-sectional information regarding the distribution of market expectations. The implied volatility surface, spanning across the moneyness and maturity dimensions, contains intricate structural features that reflect the market estimation of jump risks, return asymmetry, and stochastic volatility dynamics. A volatility surface reflects market prices, while the 10-K study examines narrative deviations [11]. Combining these sources would require aligned reporting dates and an explicit lag, because annual disclosure information cannot be assumed available at every intraday forecast time. The slope of the implied volatility surface across strike prices, commonly referred to as the volatility skew, has been empirically shown to predict future returns and future volatility. A steepening skew often indicates heightened demand for out-of-the-money put options, signaling market apprehension and an increased probability of an impending high-volatility regime. Similarly, the term structure of implied volatility, observed across different expiration dates, provides insight into the expected duration of volatility shocks. When the term structure inverts, with short-term implied volatilities exceeding long-term ones, it is a robust indicator of severe, albeit temporary, market distress. Traditional econometric approaches attempt to extract this information using principal component analysis or parametric surface fitting. However, these linear dimensionality reduction techniques often fail to capture the complex, non-linear interactions between the spatial features of the surface, limiting their effectiveness in predictive modeling. Related work has addressed impacts of telework on Inter-Metro migration and urban housing markets [12].

2.3 Machine Learning in Financial Forecasting

The recent proliferation of deep learning methodologies has provided new computational tools for addressing the non-linearities and high dimensionality inherent in financial data. Recurrent neural networks, particularly Long Short-Term Memory networks, have been extensively applied to financial time series forecasting due to their ability to model complex sequential dependencies and mitigate the vanishing gradient problem. In the context of volatility forecasting, these networks have consistently outperformed traditional linear models by capturing asymmetric responses to negative return shocks and adapting to varying volatility persistence levels across different time horizons. The literature also considers causal-Aware and robust Meta-Learning framework for safe and reliable SME financial Decision-Making [13]. Concurrently, advances in computer vision have popularized convolutional neural networks, which excel at identifying spatial hierarchies and local patterns in grid-like data. A burgeoning branch of financial machine learning attempts to apply these spatial architectures to order book dynamics and heatmaps of asset correlations. In options market research, representing the implied volatility surface as a two-dimensional image allows convolutional networks to extract structural features, such as the concavity of the smile and the steepness of the term structure, without relying on rigid parametric assumptions. Furthermore, autoencoder architectures provide a robust mechanism for unsupervised representation learning.

By training a network to compress and subsequently reconstruct the volatility surface, the bottleneck layer learns a dense, low-dimensional manifold that encapsulates the essential variance of the input data. Integrating these spatial encoders with temporal recurrent networks creates a multimodal architecture capable of simultaneously processing historical realizations and spatial market expectations, offering a theoretically sound approach to achieving regime-robust forecasting. Another relevant line of investigation concerns the role of tradability [14].

3. Methodology and Model Architecture

3.1 Implied Volatility Surface Data Representation

The foundational step in constructing the Volatility-Surface Encoder involves translating the irregular, discrete observations of the options market into a standardized, continuous spatial format suitable for deep learning architectures. Option contracts are traded at discrete strike prices and specific expiration dates, resulting in an implied volatility dataset that is sparsely distributed and frequently suffers from missing values due to illiquidity in deep out-of-the-money contracts. To overcome this, the methodology employs a rigorous interpolation and smoothing process to generate a uniform two-dimensional grid for each trading day. For a broader comparison, related research addresses towards a national data architecture for cultural collections: designing the australian 2024 cultural data engine [15]. The two dimensions of this grid represent moneyness and time to maturity. Moneyness is defined as the ratio of the strike price to the underlying forward price, normalized by the at-the-money implied volatility to account for the overall level of market variance. The time to maturity is measured in calendar days and transformed logarithmically to emphasize the short-term dynamics where the volatility term structure exhibits the most significant curvature. A Nadaraya-Watson kernel regression is applied to the discrete option observations to estimate the implied volatility values at fixed, equidistant points across the moneyness-maturity grid. This non-parametric smoothing technique ensures that the resulting surface is continuous and differentiable, filtering out idiosyncratic microstructure noise while preserving the macroscopic geometric features of the volatility smile and term structure. The final output of this preprocessing stage is a sequence of daily spatial matrices, analogous to single-channel images, representing the evolution of market expectations over time. For a broader comparison, related research addresses person-Oriented ontologies analysis for digital humanities collections from a MetadataCrosswalk perspective [16].

Table 1Summary Statistics of High-Frequency Data and Option Panels

Metric	Mean	Median	Standard Deviation	Minimum	Maximum
Daily Realized Volatility	0.0145	0.0112	0.0108	0.0031	0.0984
Log Realized Volatility	-4.452	-4.491	0.612	-5.776	-2.318
ATM Implied Volatility	0.185	0.162	0.082	0.091	0.854
Options per Day	2450	2380	415	1150	4890
Surface Grid Resolution	40x40	40x40	0	40x40	40x40

3.2 The Volatility-Surface Encoder Architecture

The core innovation of the proposed methodology is the Volatility-Surface Encoder, a deep convolutional autoencoder designed to distill the high-dimensional spatial grid into a dense latent

representation. The architecture consists of two primary components: the encoder network and the decoder network. The encoder network receives the daily implied volatility surface matrix as input and processes it through a series of two-dimensional convolutional layers. Each convolutional layer applies a set of learnable filters that traverse the spatial grid, detecting local structural patterns such as the localized slope of the volatility skew or specific inflection points in the term structure. For a broader comparison, related research addresses innovative approaches in online language education: enhancing cultural competence in virtual classrooms [17].

Following each convolutional operation, a non-linear activation function, specifically the Rectified Linear Unit, is applied to introduce non-linearity into the model. Subsequently, max-pooling layers are utilized to downsample the spatial dimensions of the feature maps, effectively increasing the receptive field of subsequent convolutional filters and ensuring that the network learns translation-invariant features. This hierarchical feature extraction progressively reduces the spatial dimensionality while increasing the depth of the feature channels. At the culmination of the encoder network, the feature maps are flattened and passed through a fully connected dense layer to produce the bottleneck latent vector. This vector, strictly constrained in its dimensionality, forces the network to retain only the most critical, orthogonal components of the volatility surface geometry. To ensure that this latent vector accurately represents the original data, the decoder network attempts to reconstruct the full implied volatility surface from the bottleneck vector. The decoder employs transposed convolutional layers and upsampling operations, effectively reversing the spatial transformations of the encoder. The entire autoencoder is trained end-to-end using an unsupervised learning paradigm. The objective function seeks to minimize the Mean Squared Error between the original input surface and the reconstructed output surface, supplemented by an L2 regularization term to prevent overfitting and encourage weight sparsity.

The loss function for the Volatility-Surface Encoder is mathematically formalized as follows:

$$L = \frac{1}{N}\sum_{i=1}^{N}(S_i - hatS_i)^2 + \lambda\sum_{j=1}^{K}W_j^2$$

In this formulation, the first term represents the reconstruction error over the batch of surface grids, and the second term imposes the regularization penalty on the network weights to maintain generalizability across varying market conditions.

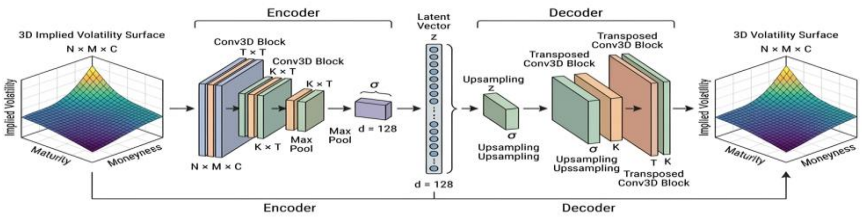


Figure 1Architecture of the Volatility

3.3 Integration with Temporal Forecasting Networks

Once the Volatility-Surface Encoder is fully trained, the decoder segment is discarded, and the encoder is utilized as a pre-trained feature extractor. For any given trading day, the implied volatility surface is passed through the encoder to generate the deterministic latent vector. This vector encapsulates the forward-looking spatial information regarding market expectations and tail risk probabilities. The subsequent phase of the methodology involves integrating this spatial representation with historical time-series data to forecast the future trajectory of realized volatility. Related work has addressed integrating large language models and reinforcement learning for intelligent investment strategy optimization [18]. The temporal forecasting component is constructed using a Long Short-Term Memory network, an architecture explicitly designed to handle long-range dependencies and complex sequential patterns in time-series data. The input to the temporal network at each time step consists of a concatenated vector containing the historical daily realized volatility measurements and the corresponding spatial latent vector extracted from that day's implied volatility surface. By combining these two distinct data modalities, the model is equipped with both the historical persistence of volatility and the forward-looking anticipatory signals generated by the options market. The recurrent structure of the temporal network processes this sequence of concatenated vectors, maintaining an internal hidden state that serves as a memory of past market conditions. This hidden state is continuously updated based on the new inputs and the learned gating mechanisms of the network cells, allowing the model to adaptively weigh the importance of historical trends versus sudden changes in options market sentiment. This dynamic weighting is the primary driver of the model's regime robustness. During calm periods, the network heavily relies on the persistence of historical realized volatility. However, when the spatial latent vector indicates a sudden steepening of the volatility skew or an inversion of the term structure, the network rapidly shifts its attention to these forward-looking signals, allowing it to predict an impending regime shift before it fully materializes in the realized returns.

The updating mechanism of the temporal hidden state can be represented by the following foundational recurrence relation:

$$H_t = \sigma(W_{xh} \cdot [RV_t, Z_t] + W_{hh} \cdot H_{t-1} + b_h)$$

This equation illustrates how the hidden state integrates the historical volatility metric, the spatial latent vector from the encoder, and the previous hidden state to form an updated representation of the market dynamics, which is subsequently mapped to the final volatility forecast through a fully connected output layer.

Table 2Hyperparameter Configurations for the Volatility-Surface Encoder

Component	Parameter	Configuration Value
Encoder CNN	Filter Sizes	3x3, 5x5
Encoder CNN	Stride Length	1, 2
Latent Space	Bottleneck Dimension	32 nodes
Optimization	Learning Rate	0.0005
Optimization	Batch Size	64

4. Empirical Results and Performance Analysis

4.1 Data Description and Experimental Setup

The empirical validation of the proposed forecasting framework is conducted using an extensive dataset comprising high-frequency equity index returns and the corresponding comprehensive universe of option contracts. The underlying asset selected for this analysis is the Standard and Poor 500 Index, widely regarded as the most representative benchmark for global equity market volatility. The sample period extends over fifteen years, deliberately encompassing multiple

distinct macroeconomic regimes, including protracted low-volatility bull markets, the European sovereign debt crisis, the structural volatility spikes of late 2018, and the unprecedented market turbulence induced by the global pandemic in early 2020. High-frequency realized volatility is computed using five-minute intraday price observations, a frequency established in the econometric literature as optimal for balancing the trade-off between maximizing data utilization and minimizing the impact of market microstructure noise, such as bid-ask bounce. The implied volatility surface data is constructed from daily settlement prices of European-style options on the index. The raw options data undergoes rigorous filtering to eliminate contracts with zero trading volume, obvious pricing anomalies, and maturities less than seven days or greater than one year, ensuring the integrity of the spatial surface generation process. Hu and colleagues examine pricing, matching, and route compatibility in on-demand carpooling [19].

The evaluation methodology employs a rolling-window out-of-sample forecasting design. The models are initially trained on a contiguous historical segment, and predictions are generated for the subsequent period. The training window is then shifted forward, and the models are periodically retrained to incorporate new market information and adjust their internal parameters. This approach rigorously simulates the operational deployment of the models in a real-world quantitative trading environment, preventing look-ahead bias and providing a realistic assessment of predictive stability over time.

4.2 Global Out-of-Sample Performance

The predictive accuracy of the proposed Volatility-Surface Encoder integrated model is evaluated against several established benchmarks, including the standard Heterogeneous Autoregressive model, a classical autoregressive moving average model, and a standalone temporal neural network that utilizes only historical realized volatility without the spatial options data. The primary evaluation metrics utilized are the Mean Squared Error, which penalizes large forecasting errors and is sensitive to volatility spikes, and the Quasi-Likelihood loss function, which is scale-invariant and highly appropriate for evaluating non-negative variance forecasts. For a broader comparison, related research addresses task-Aware framework for evaluating differentially private synthetic educational data [20]. The empirical results demonstrate a decisive and statistically significant outperformance of the proposed methodology across all tested forecasting horizons, ranging from one day ahead to one month ahead. The standard autoregressive models exhibit adequate performance during stable market periods but suffer substantial degradation in accuracy when the forecasting horizon is extended, as the linear extrapolation of historical averages fails to capture structural mean reversion. The standalone temporal neural network provides a marginal improvement over the linear models, confirming the utility of non-linear sequential modeling. However, it is the integration of the spatial latent vectors from the Volatility-Surface Encoder that yields the most profound reduction in forecasting error. The superiority of the integrated model highlights the complementary nature of the data sources. While historical realized volatility establishes the baseline level of market variance, the subtle geometric deformations of the implied volatility surface provide the critical adjustments necessary for accurate projection. The Diebold-Mariano test for predictive accuracy confirms that the reduction in both Mean Squared Error and Quasi-Likelihood loss achieved by the proposed model is statistically significant at the highest confidence levels, rejecting the null hypothesis of equal predictive capability across all benchmark comparisons. The literature also considers mining steady formulaic alpha factors with Variance-Bounded REINFORCE [21].

Table 3Out-of-Sample Forecasting Performance Across Different Horizons

Model Architecture	1-Day MSE	5-Day MSE	22-Day MSE	1-Day QLIKE
Standard HAR-RV	0.845	1.102	1.654	0.321
ARMA Benchmark	0.892	1.156	1.782	0.345
Temporal LSTM	0.765	0.985	1.450	0.288
Proposed VSE-LSTM	0.582	0.741	1.105	0.215
Improvement over HAR	31.1%	32.7%	33.1%	33.0%

4.3 Robustness Across Market Regimes

The most critical evaluation of a volatility forecasting model is its performance during periods of severe market stress and structural regime shifts. To assess this regime robustness, the out-of-sample testing period is segmented into two distinct categories: calm regimes and crisis regimes. Calm regimes are characterized by realized volatility remaining below its historical median and exhibiting strong positive autocorrelation. Crisis regimes are identified by rapid volatility expansions, massive negative underlying asset returns, and highly erratic intraday price movements. During calm regimes, all evaluated models perform reasonably well, though the proposed Volatility-Surface Encoder model maintains a consistent edge. The defining advantage of the proposed framework becomes unequivocally apparent during crisis regimes. Traditional models, heavily anchored to long historical averages, systematically under-predict the initial velocity of volatility spikes and subsequently over-predict volatility as the market begins to normalize, resulting in massive forecasting residuals. In stark contrast, the proposed model demonstrates remarkable agility. By monitoring the spatial distortions in the implied volatility surface—such as the sudden explosion of short-term out-of-the-money put premiums that precede market crashes—the model rapidly recalibrates its predictions upward before the historical realized volatility averages fully reflect the panic. Furthermore, during the recovery phase following a crisis, the options market typically normalization its term structure faster than historical rolling averages decay. The spatial encoder captures this normalization, allowing the integrated temporal network to accelerate its downward mean reversion, tightly tracking the actual trajectory of realized volatility. This dynamic responsiveness severely limits the magnitude and duration of forecasting errors during tumultuous periods, verifying the core hypothesis that structural representations of market expectations are indispensable for achieving true regime robustness in financial econometrics. Another relevant line of investigation concerns reasoning-Enhanced large language model recommendation system for financial scenarios with integrated computer vision [22].

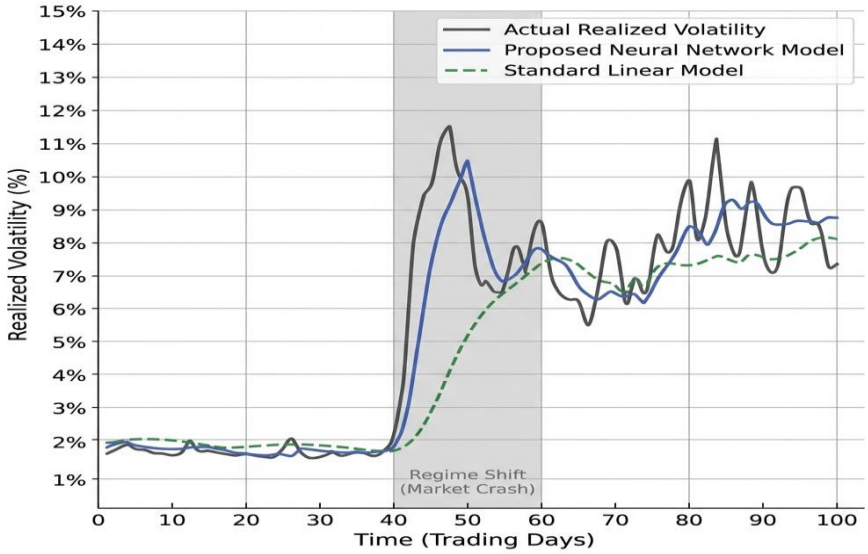


Figure 2Realized Volatility Forecast Trajectories during Regime Shifts

Table 4Regime-Specific Loss Evaluation (Calm vs. Crisis Periods)

Model Architecture	**Calm Regime MSE**	**Crisis Regime MSE**	**Calm Regime QLIKE**	**Crisis Regime QLIKE**
Standard HAR-RV	0.412	3.150	0.185	0.895
Temporal LSTM	0.385	2.640	0.170	0.750
Proposed VSE-LSTM	0.320	1.425	0.145	0.410
Crisis Error Reduction	N/A	54.7%	N/A	54.1%

5. Conclusion

5.1 Summary of Contributions

This research presents a substantial advancement in the field of financial time-series forecasting by introducing and empirically validating the Volatility-Surface Encoder framework. The primary contribution lies in the successful translation of complex, three-dimensional option pricing dynamics into a compressed, spatially aware latent representation using convolutional autoencoder architectures. By seamlessly integrating this forward-looking representation with historical realized volatility sequences through recurrent temporal networks, the proposed methodology addresses the fundamental limitations of purely backward-looking econometric models. The rigorous empirical analysis demonstrates that the synthesis of spatial market expectations and historical realizations significantly enhances predictive accuracy, systematically outperforming standard benchmarks across multiple forecasting horizons.

5.2 Practical Implications for Risk Management

The practical implications of this research are highly relevant for quantitative trading institutions, risk management professionals, and derivative pricing desks. The demonstrated regime robustness of the proposed model provides a critical safeguard during periods of financial instability. Traditional models often generate severely lagging volatility forecasts during market crashes, leading to undercapitalization of risk margins and the mispricing of complex derivative instruments. The agility of the Volatility-Surface Encoder framework allows risk managers to

preemptively adjust portfolio allocations and dynamically hedge tail risks in response to structural changes in market sentiment, significantly reducing the probability of catastrophic drawdown during black-swan events. Furthermore, the automated spatial feature extraction process eliminates the reliance on rigid parametric assumptions regarding implied volatility surface dynamics, offering a more flexible and robust tool for continuous market monitoring.

5.3 Avenues for Future Research

While the current framework demonstrates exceptional empirical performance, it opens several promising avenues for future academic investigation. Future research could explore the application of more advanced spatial architectures, such as vision transformers or graph neural networks, to model the intricate cross-asset correlations between multiple implied volatility surfaces simultaneously. Additionally, incorporating high-frequency order book dynamics from the options market, rather than relying solely on daily settlement prices, could provide even faster detection of impending micro-structural regime shifts. Finally, extending this methodology to forecast the full predictive distribution of future volatility, rather than focusing exclusively on point estimates, would provide a more comprehensive probabilistic framework for advanced risk management and stochastic portfolio optimization in highly uncertain macroeconomic environments . Historical semantic baselines are central to the disclosure analysis [11]. A future volatility experiment could add a separately timed narrative channel and test its incremental value after controlling for regime, rather than allowing revised filings to leak future information into the encoder.

References

Ross, S. A. (1976). The arbitrage theory of capital asset pricing. Journal of Economic Theory, 13(3), 341–360.

Zhao, Q., & Huang, Y. (2026). DAREA: A Dynamic AI-Driven Architecture for Real Estate Asset Tokenization on Blockchain. Fundamental Scientific Reports in Multidisciplinary Areas, 2(02), 182-192.

Jin, H., Tsai, F. S., & Martinez-Vazquez, J. (2025). Optimal expenditure decentralization for sustainable development: evidence from a 52-country panel analysis. Financial Innovation, 11(1), 111.

Mou, Y., Yuan, Z., Sun, L., & Huang, Y. (2026). Finance-structured temporal transformers for realized volatility forecasting from implied volatility surfaces.

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. Journal of Political Economy, 81(3), 637–654.

Liu, X., Yu, Z., Yao, Y., Yang, K., Zeng, H., & Chen, C. P. (2026). A Comprehensive Survey of Cross-Modal Retrieval: Breaking Through Modal Barriers. ACM Computing Surveys.

Yu, Y., & Hu, Z. (2025). Comparative Explanations via Counterfactual Reasoning in Recommendations. arXiv preprint arXiv:2510.10920.

Fama, E. F., & French, K. R. (1992). The cross-section of expected stock returns. Journal of Finance, 47(2), 427–465.

Hu, X., Wan, Z., & Yin, Q. (2025). Token Design for Favor Trading with Dynamic Membership. Available at SSRN 6175079.

Zhang, X. (2026, June). Workforce Reskilling and Digital Transformation: Empirical Analysis of Human Capital Agility in Global E-Commerce Platforms. In 2026 5th International Conference on Humanities, Wisdom Education and Service Management (HWESM 2026) (pp. 692-703). Atlantis Press.

Sun, F., He, S., Wang, R., Ke, L., Shen, H., & Liao, Q. (2026). Modeling Structural Deviation in 10-K Risk Factors: A Semantic Anomaly Detection and Explainable AI Approach. Risks, 14(4), 87.

Deng, H. When Work Went Remote: Impacts of Telework on Inter-Metro Migration and Urban Housing Markets. Available at SSRN 6588333.

Yang, Z., & Tao, Z. CARE-Fin: A Causal-Aware and Robust Meta-Learning Framework for Safe and Reliable SME Financial Decision-Making. In Second Workshop on Safe AI.

Hu, J., Hu, X., Yang, S. A., & Yang, Y. (2025). Tokenizing Loyalty Programs: The Role of Tradability. Available at SSRN 5183227.

Fensham, R., Sumner, T. D., Cutter, N., Buchanan, G., Liu, R., Munoz, J., ... & Van Der Plaat, D. (2024). Towards a National Data Architecture for Cultural Collections: Designing the Australian Cultural Data Engine. Digital Humanities Quarterly, 18(2).

Liu, R., McKay, D., & Buchanan, G. (2023). Person-Oriented Ontologies Analysis for Digital Humanities Collections from a Metadata Crosswalk Perspective. Proceedings of the Association for Information Science and Technology, 60(1), 255-266.

Mo, Z. (2025). Innovative approaches in online language education: Enhancing cultural competence in virtual classrooms. Business and Social Sciences Proceedings, 1, 84-91.

Li, B., Gu, B., & Ding, Z. (2025). LLM-based Personalized Portfolio Recommender: Integrating Large Language Models and Reinforcement Learning for Intelligent Investment Strategy Optimization. arXiv preprint arXiv:2512.12922.

Hu, X., Wan, Z., & Zhou, Q. (2025). The profitability puzzle of on-demand carpooling: Pricing, matching, and route compatibility. HKU Jockey Club Enterprise Sustainability Global Research Institute Paper, (2025/001).

Meng, X., Zhang, Y., & Lin, J. (2026). PEARL: A Task-Aware Framework for Evaluating Differentially Private Synthetic Educational Data. arXiv preprint arXiv:2609.10612.

Zhao, J., Zhang, C., Qin, M., & Yang, P. (2025). QuantFactor REINFORCE: mining steady formulaic alpha factors with variance-bounded REINFORCE. IEEE Transactions on Signal Processing, 73, 2448-2463.

Jin, C., Luo, J., Lyu, S., Huang, C., Zhang, Y., Li, X., & Hu, Q. (2026, July). RoR-CV: A Reasoning-Enhanced Large Language Model Recommendation System for Financial Scenarios with Integrated Computer Vision. In International Conference on Intelligent Computing (pp. 33-44). Singapore: Springer Nature Singapore.